

---

## Search Combinators

Tom Schrijvers · Guido Tack · Pieter Wuille ·  
Horst Samulowitz · Peter J. Stuckey

**Abstract** The ability to model search in a constraint solver can be an essential asset for solving combinatorial problems. However, existing infrastructure for defining search heuristics is often inadequate. Either modeling capabilities are extremely limited or users are faced with a general-purpose programming language whose features are not tailored towards writing search heuristics. As a result, major improvements in performance may remain unexplored.

This article introduces *search combinators*, a lightweight and solver-independent method that bridges the gap between a conceptually simple modeling language for search (high-level, functional and naturally compositional) and an efficient implementation (low-level, imperative and highly non-modular). By allowing the user to define application-tailored search strategies from a small set of primitives, search combinators effectively provide a rich *domain-specific language* (DSL) for modeling search to the user. Remarkably, this DSL comes at a low implementation cost to the developer of a constraint solver.

The article discusses two modular implementation approaches and shows, by empirical evaluation, that search combinators can be implemented without overhead compared to a native, direct implementation in a constraint solver.

---

Tom Schrijvers · Pieter Wuille  
Universiteit Gent, Belgium  
E-mail: {tom.schrijvers,pieter.wuille}@ugent.be

Guido Tack  
Monash University, Victoria, Australia  
E-mail: guido.tack@monash.edu

Pieter Wuille  
Katholieke Universiteit Leuven, Belgium  
E-mail: pieter.wuille@cs.kuleuven.be

Horst Samulowitz  
IBM Research, New York, USA  
E-mail: samulowitz@us.ibm.com

Peter J. Stuckey  
National ICT Australia (NICTA) and University of Melbourne, Victoria, Australia  
E-mail: pjs@cs.mu.oz.au

## 1 Introduction

Search heuristics often make all the difference between effectively solving a combinatorial problem and utter failure. Heuristics make a search algorithm efficient for a variety of reasons, e.g., incorporation of domain knowledge, or randomization to avoid heavy-tailed runtimes. Hence, the ability to swiftly design search heuristics that are tailored towards a problem domain is essential for performance. This article introduces search combinators, a versatile, modular, and efficiently implementable language for expressing search heuristics.

### 1.1 Status Quo

In CP, much attention has been devoted to facilitating the modeling of combinatorial problems. A range of high-level modeling languages, such as Zinc [13], OPL [29] or Comet [27], enable quick development and exploration of problem models. However, we see very little support on the side of formulating accompanying search heuristics. Most languages and systems, e.g. MiniZinc [14], Comet [27], Gecode [23], or ECLiPSe [7] provide a small set of predefined heuristics “off the shelf”. Some systems also support user-defined search based on a general-purpose programming language (e.g., all of the above systems except MiniZinc). The former is clearly too confining, while the latter leaves to be desired in terms of productivity, since implementing a search heuristic quickly becomes a non-negligible effort. This also explains why the set of predefined heuristics is typically small: it takes a lot of time for CP system developers to implement heuristics, too – time they would much rather spend otherwise improving their system.

### 1.2 Contributions

In this article we show how to resolve this stand-off between solver developers and users, by introducing a domain-specific modular search language based on combinators, as well as a modular, extensible implementation architecture.

**For the user,** we provide a modeling language for expressing complex search heuristics based on an (extensible) set of primitive combinators. Even if the users are only provided with a small set of combinators, they can already express a vast range of combinations. Moreover, using combinators to program application-tailored search is vastly more productive than resorting to a general-purpose language.

**For the system developer,** we show how to design and implement modular combinators. The modularity of the language thus carries over directly to modularity of the implementation. Developers do not have to cater explicitly for all possible combinator combinations. Small implementation efforts result in providing the user with a lot of expressive power. Moreover, the cost of adding one more combinator is small, yet the return in terms of additional expressiveness can be quite large.

The technical challenge is to bridge the gap between a conceptually simple search language and an efficient implementation, which is typically low-level, imperative and highly non-modular. This is where existing approaches fail; they restrict the expressiveness of their search specification language to face up to implementation limitations, or they raise errors when the user strays out of the implemented subset.

The contribution is therefore the novel design of an expressive, high-level, compositional search language with an equally modular, extensible, and efficient implementation architecture.

### 1.3 Approach

We overcome the modularity challenge by implementing the primitives of our search language as *mixin* components [4]. As in Aspect-Oriented Programming [10], mixin components neatly encapsulate the *cross-cutting behavior* of primitive search concepts, which are highly entangled in conventional approaches. Cross-cutting means that a mixin component can interfere with the behavior of its sub-components (in this case, sub-searches). The combination of encapsulation *and* cross-cutting behavior is essential for systematic reuse of search combinators. Without this degree of modularity, minor modifications require rewriting from scratch.

An added advantage of mixin components is extensibility. We can add new features to the language by adding more mixin components. The cost of adding such a new component is small, because it does not require changes to the existing ones. Moreover, experimental evaluation bears out that this modular approach has no significant overhead compared to the traditional monolithic approach. Finally, our approach is solver-independent and therefore makes search combinators a potential standard for designing search.

### 1.4 Plan of the Article

The rest of the article is structured as follows. The next section defines the high-level search language in terms of basic heuristics and combinators. Sect. 3 shows how the modular language is mapped to a *modular design* of the combinator implementations. Sect. 4 presents two concrete implementation approaches for combinators and gives an overview of how we integrate search combinators into the MiniZinc toolchain. Sect. 5 verifies that combinators can be implemented with low overhead. Finally, Sect. 6 discusses related approaches, and Sect. 7 concludes the article.

### 1.5 Note to Reviewers

This article is an extended version of a paper [21] that appeared in the proceedings of the 17th International Conference on Principles and Practice of Constraint Programming (CP) 2011. That paper further developed the ideas laid out in our earlier paper [19], which was presented at ModRef 2010.

Compared to the CP'11 conference version, this article features a completely rewritten and more detailed introduction of the combinator language, both from the high-level (Sect. 2) and the implementation-level (Sect. 3) point of view. We have completed the integration into MiniZinc and present a full toolchain in Sect. 4.3. In addition, the discussion of related work (Sect. 6) has been extended with much more detail.

$s$	$::=$	prune		ifthenelse( $c, s_1, s_2$ )
		prunes the node		perform $s_1$ until $c$ is false, then perform $s_2$
		base_search( $vars, var-select, domain-split$ )		and( $[s_1, s_2, \dots, s_n]$ )
		label		perform $s_1$ , on success $s_2$ otherwise fail, ...
		let( $v, e, s$ )		or( $[s_1, s_2, \dots, s_n]$ )
		introduce new variable $v$ with		perform $s_1$ , on termination start $s_2$ , ...
		initial value $e$ , then perform $s$		portfolio( $[s_1, s_2, \dots, s_n]$ )
		assign( $v, e$ )		perform $s_1$ , if not exhaustive start $s_2$ , ...
		assign $e$ to variable $v$ and succeed		restart( $c, s$ )
		post( $c, s$ )		restart $s$ as long as $c$ holds
		post constraint $c$ at every node during $s$		

**Fig. 1:** Catalog of primitive search heuristics and combinators

## 2 High-Level Search Language

This section introduces the syntax of our high-level search language and illustrates its expressive power and modularity by means of examples. The rest of the article then presents an architecture that maps the modularity of the language down to the implementation level.

The search language is used to define a *search heuristic*, which a *search engine* applies to each node of the search tree. For each node, the heuristic determines whether to continue search by creating child nodes, or to prune the tree at that node. The queuing strategy, i.e., the strategy by which new nodes are selected for further search (such as depth-first traversal), is determined separately by the search engine, it is thus orthogonal to the search language. The search language features a number of primitives, listed in the catalog of Fig. 1. These are the building blocks in terms of which more complex heuristics can be defined, and they can be grouped into *basic heuristics* (base\_search and prune), *combinators* (ifthenelse, and, or, portfolio, and restart), and *state management* (let, assign, post). This section introduces the three groups of primitives in turn.

We emphasize that this catalog is open-ended; we will see that the language implementation explicitly supports adding new primitives.

The concrete syntax we chose for presentation uses simple nested terms, which makes it compatible with the *annotation* language of MiniZinc [14]. Sect. 4.3 discusses our implementation of MiniZinc with combinator support. However, other concrete syntax forms are easily supported (e.g., we support C++ and Haskell).

### 2.1 Basic Heuristics

Let us first discuss the two basic primitives, base\_search and prune.

**base\_search.** The most widely used method for specifying a basic heuristic for a constraint problem is to define it in terms of a *variable selection* strategy which picks the next variable to constrain, and a *domain splitting* strategy which splits the set of possible values of the selected variable into two (or more) disjoint sets. Common variable selection strategies are:

- firstfail: select the variable with the smallest current domain,
- smallest: select the variable which can take the smallest possible value,
- domwdeg [2]: select the variable with smallest ratio of size of current domain and number of failures the variable has been involved in, and
- impact [18]: select the variable that will (based on past experience) reduce the raw search space of the problem the most.

Common domain splitting strategies are:

- **min**: set the variable to its minimum value or greater than its minimum,
- **max**: set the variable to its maximum value or less than its maximum,
- **median**: set the variable to its median value, or not equal to this value, and
- **split**: constrain the variable to the lower half of its range of possible values, or its upper half.

The CP community has spent a considerable amount of work on defining and exploring the above and many other variable selection and domain splitting heuristics. The provision of a flexible language for defining new basic searches is an interesting problem in its own right, but in this article we concentrate on search combinators that combine and modify basic searches.

To this end, our search language provides the primitive `base_search(vars, var-select, domain-split)`, which specifies a systematic search. If any of the variables *vars* are still not fixed at the current node, it creates child nodes according to *var-select* and *domain-split* as variable selection and domain splitting strategies respectively.

Note that `base_search` is a CP-specific primitive; other kinds of solvers provide their own search primitives. The rest of the search language is essentially solver-independent. While the solver provides few basic heuristics, the search language adds great expressive power by allowing these to be combined arbitrarily using combinators.

**prune.** The second basic primitive, `prune`, simply cuts the search tree below the current node. Obviously, this primitive is useless on its own, but we will see shortly how `prune` can be used together with combinators.

## 2.2 Combinators


The expressive power of the search language relies on combinators, which combine search heuristics (which can be basic or themselves constructed using combinators) into more complex heuristics.

**and/or.** Probably the most widely used combination of heuristics is *sequential composition*. For instance, it is often useful to first label one set of problem variables before starting to label a second set. The following heuristic uses the `and` combinator to first label all the *xs* variables using a first-fail strategy, followed by the *ys* variables with a different strategy:

```
and([base_search(xs, firstfail, min),
     base_search(ys, smallest, max)])
```

As you can see in Fig. 1, the `and` combinator accepts a list of searches  $s_1, \dots, s_n$ , and performs their and-sequential composition. And-sequential means, intuitively, that solutions are found by performing *all* the sub-searches sequentially down one branch of the search tree, as illustrated in Fig. 2.1.

The dual combinator, `or([s1, ..., sn])`, performs a disjunctive combination of its sub-searches – a solution is found using *any* of the sub-searches (Fig. 2.2), trying them in the given order.



**Fig. 2:** Primitive combinators

**Statistics and ifthenelse.** The ifthenelse combinator is centered around a conditional expression  $c$ . As long as  $c$  is true for the current node, the sub-search  $s_1$  is used. Once  $c$  is false,  $s_2$  is used for the complete subtree below the current node (see Fig. 2.3).

We do not specify the *expression language* for conditions in detail, we simply assume that it comprises the typical arithmetic and comparison operators and literals that require no further explanation. It is notable though that the language can refer to the constraint variables and parameters of the underlying model. Additionally, a condition may refer to one or more *statistics* variables. Such statistics are collected for the duration of a subsearch until the condition is met. For instance `ifthenelse(depth < 10,  $s_1$ ,  $s_2$ )` maintains the search depth statistic during subsearch  $s_1$ . At depth 10, the ifthenelse combinator switches to subsearch  $s_2$ .

We distinguish two forms of statistics: *Local statistics* such as depth and discrepancies express properties of individual nodes. *Global statistics* such as number of explored nodes, encountered failures, solution, and time are computed for entire search trees.

It is worthwhile to mention that developers (and advanced users) can also define their own statistics, just like combinators, to complement any predefined ones. In fact, Sect. 3 will show that statistics can be implemented as a *subtype* of combinators that can be queried for the statistic's value.

**Abstraction.** Our search language draws its expressive power from the combination of primitive heuristics using combinators. An important aspect of the search language is *abstraction*: the ability to create new combinators by effectively defining macros in terms of existing combinators.

For example, we can define the limiting combinator `limit( $c$ ,  $s$ )` to perform  $s$  while condition  $c$  is satisfied, and otherwise cut the search tree using `prune`:

$\text{limit}(c, s) \equiv \text{ifthenelse}(c, s, \text{prune})$

The well-known  $\text{once}(s)$  combinator is a special case of the limiting combinator where the number of solutions is less than one. This is simply achieved by maintaining and accessing the solutions statistic:

$\text{once}(s) \equiv \text{limit}(\text{solutions} < 1, s)$

**Exhaustiveness and portfolio/restart.** The behavior of the final two combinators, *portfolio* and *restart*, depends on whether their sub-search was *exhaustive*. Exhaustiveness simply means that the search has explored the entire subtree without ever invoking the *prune* primitive.

The  $\text{portfolio}([s_1, \dots, s_n])$  combinator performs  $s_1$  until it has explored the whole subtree. If  $s_1$  was exhaustive, i.e., if it did not call *prune* during the exploration of the subtree, the search is finished. Otherwise, it continues with  $\text{portfolio}([s_2, \dots, s_n])$ . This is illustrated in Fig. 2.4, where the subtree of  $s_1$  represents a non-exhaustive search,  $s_2$  is exhaustive and therefore  $s_3$  is never invoked.

An example for the use of *portfolio* is the  $\text{hotstart}(c, s_1, c_2)$  combinator. It performs search heuristic  $s_1$  while condition  $c$  holds to initialize global parameters for a second search  $s_2$ . This heuristic can for example be used to initialize the widely applied *Impact* heuristic [18]. Note that we assume here that the parameters to be initialized are maintained by the underlying solver, so we omit an explicit reference to them.

$\text{hotstart}(c, s_1, s_2) \equiv \text{portfolio}([\text{limit}(c, s_1), s_2])$

The  $\text{restart}(c, s)$  combinator repeatedly runs  $s$  in full. If  $s$  was not exhaustive, it is restarted, until condition  $c$  no longer holds. Fig. 2.5 shows the two cases, on the left terminating with an exhaustive search  $s$ , on the right terminating because  $c$  is no longer true.

The following implements random restarts, where search is stopped after 1000 failures and restarted with a random strategy:

$\text{restart}(\text{true}, \text{limit}(\text{failures} < 1000, \text{base\_search}(xs, \text{randomvar}, \text{randomval})))$

Clearly, this strategy has a flaw: If it takes more than 1000 failures to find the solution, the search will never finish. We will shortly see how to fix this by introducing user-defined search variables.

The *prune* primitive is the only source of non-exhaustiveness. Combinators propagate exhaustiveness in the obvious way:

- $\text{and}([s_1, \dots, s_n])$  is exhaustive if all  $s_i$  are
- $\text{or}([s_1, \dots, s_n])$  is exhaustive if all  $s_i$  are
- $\text{portfolio}([s_1, \dots, s_n])$  is exhaustive if one  $s_i$  is
- $\text{restart}(c, s)$  is exhaustive if the last iteration is exhaustive
- $\text{ifthenelse}(c, s_1, s_2)$  is exhaustive if  $s_1$  and  $s_2$  are

## 2.3 State Access and Manipulation

The remaining three primitives, *let*, *assign*, and *post*, are used to access and manipulate the state of the search:

- `let( $v, e, s$ )` introduces a new search variable  $v$  with initial value of the expression  $e$  and visible in the search  $s$ , then continues with  $s$ . Note that search variables are distinct from the decision variables of the model.
- `assign( $v, e$ )`: assigns the value of the expression  $e$  to search variable  $v$  and succeeds.
- `post( $c, s$ )`: provides access to the underlying constraint solver, posting a constraint  $c$  at every node during  $s$ . If  $s$  is omitted, it posts the constraint and immediately succeeds.

These primitives add a great deal of expressivity to the language, as the following examples demonstrate.

*Random restarts:* Let us reconsider the example using random restarts from the previous section, which suffered from incompleteness because it only ever explored 1000 failures. A standard way to make this strategy complete is to increase the limit geometrically with each iteration:

```
geom_restart(s)  $\equiv$  let(maxfails, 100,
                      restart(true, portfolio([limit(failures < maxfails, s),
                                                assign(maxfails, maxfails * 1.5),
                                                prune])))
```

The search initializes the search variable *maxfails* to 100, and then calls search  $s$  with *maxfails* as the limit. If the search is exhaustive, both the portfolio and the restart combinators are finished. If the search is not exhaustive, the limit is multiplied by 1.5, and the search starts over. Note that `assign` succeeds, so we need to call `prune` afterwards in order to propagate the non-exhaustiveness of  $s$  to the `restart` combinator.

*Branch-and-bound:* A slightly more advanced example is the branch-and-bound optimization strategy:

```
bab(obj, s)  $\equiv$  let(best,  $\infty$ , post(obj < best, and([s, assign(best, obj)])))
```

It introduces a variable *best* that initially takes value  $\infty$  (for minimization). In every node, it posts a constraint to bound the objective variable by *best*. Whenever a new solution is found, the bound is updated accordingly using `assign`.

The `bab` example demonstrates how search variables (like *best*) and model variables (like *obj*) can be mixed in expressions. This makes it possible to remember the state of the search between invocations of a heuristic. All of the following combinators make use of this feature.

*Restarting branch-and-bound:* This is a twist on regular branch-and-bound that restarts whenever a solution is found.

```
restart_bab(obj, s)  $\equiv$  let(best,  $\infty$ , restart(true, and([post(obj < best), once(s),
                                                         assign(best, obj)])))
```

*Radiotherapy treatment planning:* The following search heuristic can be used to solve radiotherapy treatment planning problems [1]. The heuristic minimizes a variable  $k$  using branch-and-bound (`bab`), first searching the variables  $N$ , and then verifying the solution by partitioning the problem along the  $row_i$  variables for each row  $i$  one at a time (expressed as a MiniZinc array comprehension). Failure on one row must be caused by the search on the variables in  $N$ , and consequently search never backtracks into other rows.



This behavior is similar to the `once` combinator defined above. However, when a single solution is found, the search should be considered exhaustive. We therefore need an exhaustive variant of `once`, which can be implemented by replacing `prune` with `post(false)`:

$$\text{exh\_once}(s) \equiv \text{ifthenelse}(\text{solutions} < 1, s, \text{post}(\text{false}))$$

This allows us to express the entire search strategy for radiotherapy treatment planning:

$$\text{bab}(k, \text{and}([\text{base\_search}(N, \dots)] ++ [\text{exh\_once}(\text{base\_search}(\text{row}_i, \dots)) \mid i \text{ in } 1..n]))$$

*For:* The `for` loop construct ( $v \in [l, u]$ ) can be defined as:

$$\text{for}(v, l, u, s) \equiv \text{let}(v, l, \text{restart}(v \leq u, \text{portfolio}([s, \text{and}([\text{assign}(v, v + 1), \text{prune}]])]))$$

It simply runs  $u - l + 1$  times the search  $s$ , which of course is only sensible if  $s$  makes use of side effects or the loop variable  $v$ . As in the `geom_restart` combinator above, `prune` propagates the non-exhaustiveness of  $s$  to the `restart` combinator.

*Limited discrepancy search* [9] with an upper limit of  $l$  discrepancies for an underlying search  $s$ .

$$\text{lds}(l, s) \equiv \text{for}(n, 0, l, \text{limit}(\text{discrepancies} \leq n, s))$$

The `for` construct iterates the maximum number of discrepancies  $n$  from 0 to  $l$ , while `limit` executes  $s$  as long as the number of discrepancies is smaller than  $n$ . The search makes use of the discrepancies statistic that is maintained by the search infrastructure. The original LDS [9] visits the nodes in a specific order. The search described here visits the same nodes in the same order of discrepancies, but possibly in a different individual order – as this is determined by the global queuing strategy.

The following is a combination of branch-and-bound and limited discrepancy search for solving job shop scheduling problems, as described in [9]. The heuristic searches the Boolean variables `prec`, which determine the order of all pairs of tasks on the same machine. As the order completely determines the schedule, we then fix the start times using `exh_once`.

$$\text{bab}(\text{makespan}, \text{lds}(\infty, \text{and}([\text{base\_search}(\text{prec}, \dots), \text{exh\_once}(\text{base\_search}(\text{start}, \dots)]))))$$

Fully expanded, this heuristic consists of 17 combinators and is 11 combinators deep.

*Iterative deepening* [11] for an underlying search  $s$  is a particular instance of the more general pattern of restarting with an updated bound, which we have already seen in the `geom_restart` example. Here, we generalize this idea:

$$\begin{aligned} \text{id}(s) &\equiv \text{ir}(\text{depth}, 0, +, 1, \infty, s) \\ \text{ir}(p, l, \oplus, i, u, s) &\equiv \text{let}(n, l, \text{restart}(n \leq u, \text{and}([\text{assign}(n, n \oplus i), \text{limit}(p \leq n, s)]))) \end{aligned}$$

With `let`, bound  $n$  is initialized to  $l$ . Search  $s$  is pruned when statistic  $p$  exceeds  $n$ , but iteratively restarted by `restart` with  $n$  updated to  $n \oplus i$ . The repetition stops when  $n$  exceeds  $u$  or when  $s$  has been fully explored. The bound increases geometrically, if we supply  $*$  for  $\oplus$ , as in the `restart_flip` heuristic:

$$\text{restart\_flip}(p, l, i, u, s_1, s_2) \equiv \text{let}(\text{flip}, 1, \text{ir}(p, l, *, i, u, \text{and}([\text{assign}(\text{flip}, 1 - \text{flip}), \text{ifthenelse}(\text{flip} = 1, s_1, s_2)])))$$

The `restart_flip` search alternates between two search heuristics  $s_1$  and  $s_2$ . Using this as its default strategy in the *free search* category, the lazy clause generation solver *Chuffed* scored most points in the 2010 and 2011 MiniZinc Challenges.<sup>1</sup>

*Probe search*: Try out two searches  $s_1$  and  $s_2$  to a limited extent defined by condition  $c$ . Then, for the remainder, use the search that resulted in the best solution so far.

```
probe( $c, obj, s_1, s_2$ )  $\equiv$  let( $best_1, \infty$ , let( $best_2, \infty$ ,
    portfolio([ limit( $c$ , and( $[s_1, assign(best_1, obj)]$ ))
               limit( $c$ , and( $[s_2, assign(best_2, obj)]$ ))
               ifthenelse( $best_1 \leq best_2, s_1, s_2$ )]))
```

*Dichotomic search* [24] solves an optimization problem by repeatedly partitioning the interval in which the possible optimal solution can lie. It can be implemented by restarting as long the lower bound has not met the upper bound (line 2), computing the middle (line 3), and then using an or combinator to try the lower half (line 5). If it succeeds,  $obj - 1$  is the new upper bound, otherwise, the lower bound is increased (line 6).

```
dicho( $s, obj, lb, ub$ )  $\equiv$  let( $l, lb$ , let( $u, ub$ ,
    restart( $l < u$ ,
    let( $h, l + \lceil (u - l) / 2 \rceil$ ,
    once(or([
        and([post( $l \leq obj \leq h$ ),  $s$ , assign( $u, obj - 1$ )]),
        and([assign( $l, h + 1$ ), prune]))))
    ))))
```

### 3 Modular Combinator Design

The previous section caters for the user's needs, presenting a high-level modular syntax for our combinator-based search language. To cater for the system developer's needs, this section goes beyond modularity of syntax, introducing modularity of *design*.

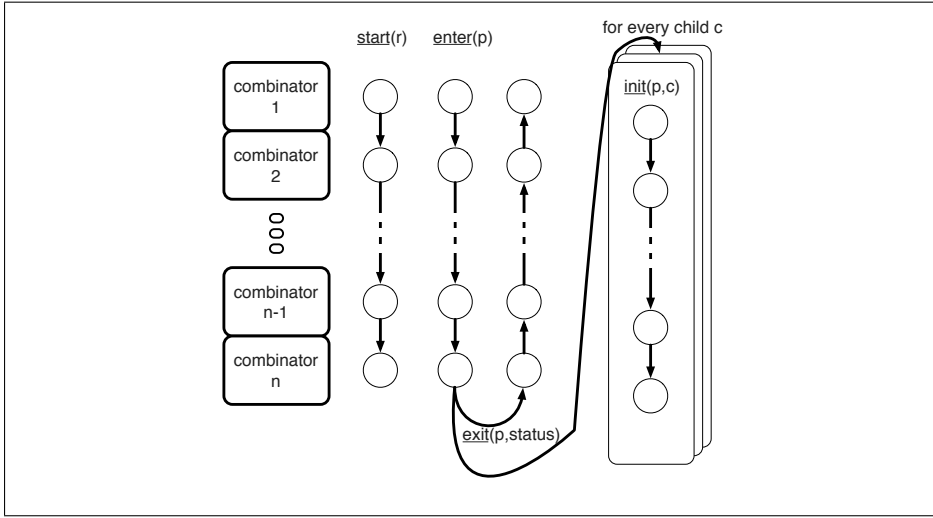
*Modularity of design* is the one property that makes our approach practical. Each combinator corresponds to a separate module that has a meaning and an implementation independent of the other combinators. This enables us to actually realize the search specifications defined by modular syntax.

Modularity of design also enables growing a system from a small set of combinators (e.g., those listed in Fig. 1), gradually adding more as the need arises. Advanced users can complement the system's generic combinators with a few application-specific ones.

*Solver independence* is another notable property of our approach. While a few combinators access solver-specific functionality (e.g., `base_search` and `post`), the approach as such and most combinators listed in Fig. 1 are in fact generic (solver- and even CP-independent); their design and implementation is reusable.

The solver-independence of our approach is reflected in the minimal interface that solvers must implement. This interface consists of an abstract type `State` which represents a state

<sup>1</sup> <http://www.g12.csse.unimelb.edu.au/minizinc/challenge2011/>



**Fig. 3:** The modular message protocol

of the solver (e.g., the variable domains and accumulated constraint propagators) which supports copying. Truly no more is needed for the approach or all of the primitive combinators in Fig. 1, except for `base_search` and `post` which require CP-aware operations for querying variable domains, solver status and posting constraints, and possibly interacting with statistics maintained by the solver. Note that there need not be a 1-to-1 correspondence between an implementation of the abstract `State` type and the solver’s actual state representation; e.g., for solvers based on trailing, recomputation techniques [15] can be used. We have implementations of the interface based on both copying and trailing.

In the following we explain our design in detail by means of code implementations of most of the primitive combinators we have covered in the previous section.

### 3.1 The Message Protocol

To obtain a modular design of search combinators we step away from the idea that the behavior of a search combinator, like the `and` combinator, forms an indivisible whole; this leaves no room for interaction. The key insight here is that we must identify finer-grained steps, defining how different combinators interact at each node in the search tree. Interleaving these finer-grained steps of different combinators in an appropriate manner yields the composite behavior of the overall search heuristic, where each combinator is able to cross-cut the others’ behavior.

Considering the diversity of combinators and the fact that not all units of behavior are explicitly present in all of them, designing this protocol of interaction is non-trivial. It requires studying the intended behavior and interaction of combinators to isolate the fine-grained units of behavior and the manner of interaction. The contribution of this section is an elegant and conceptually uniform design that is powerful enough to express all the combinators presented in this article.

We present this design in the form of a *message protocol*. The protocol specifies a set of messages (i.e., an interface with one procedure for each fine-grained step) that have to be

implemented by all combinators. In pseudo-code, this protocol for combinators consists of four different messages:

```
protocol combinator
  start (rootNode);
  enter (currentNode);
  exit (currentNode, status);
  init (parentNode, childNode);
```

All of the message signatures specify one or two search tree *nodes* as parameters. Each such node keeps track of a solver *State* and the information associated by combinators to that *State*. We observe three different access patterns of nodes:

1. In keeping with the solver independence stipulated above, we will see that most combinators only query and update their associated information and do not access the underlying solver *State* at all.
2. Restarting-based combinators, like `restart` and `portfolio`, copy nodes. This means copying the underlying solver *State* and all associated information.
3. Finally, selected solver-specific combinators like `base_search` do perform solver-specific operations on the underlying *State*, like querying variable domains and posting constraints.

In addition to the message signatures, the protocol also stipulates in what order the messages are sent among the combinators (see Fig. 3). While in general a combinator composition is tree-shaped, the processing of any single search tree node  $p$  only involves a stack of combinators. For example, given `or([and([s1, s2]), and([s3, s4])])`, either  $s_1, s_2$  or  $s_3, s_4$  are *active* for  $p$ . The picture shows this stack of active combinators on the left. Every combinator in the stack has both a *super*-combinator above and a *sub*-combinator below, except for the *top* and the *bottom* combinators. The bottom is always a basic heuristic (`base_search`, `prune`, `assign`, or `post`). The important aspect to take away from the picture is the direction of the four different messages, either top-down or bottom-up.

### 3.2 Basic Setup

Before we delve into the interesting search combinators, we first present an example implementation of the basic setup consisting of a base search (`base_search`) and a search engine (`dfs`). This allows us to express overall search specifications of the form:

```
dfs (base_search (vars, var-select, domain-split)).
```

**Base Search.** We do not provide full details on a `base_search` combinator, as it is not the focus of this article. However, we will point out the aspects relevant to our protocol.

In the `enter` message, the node's solver state is propagated. Subsequently, the condition `isLeaf(c, vars)` checks whether the solver state is unsatisfiable or there are no more variables to assign. If either is the case, the exit status (respectively **failure** or **success**) is sent to the `parent` combinator. For now, the `parent` combinator is just the search engine, but later we will see how other combinators can be inserted between the search engine and the base search.

If neither is the case, the search branches depending on the variable selection and domain splitting strategies. This involves creating a child node for each branch, determining the variable and value for that child and posting the assignment to the child's state. Then, the

**top** combinator (i.e., the engine) is asked to initialize the child node. Finally the child node is pushed onto the search queue.

```

combinator base_search(vars,var-select,domain-select)
  enter(c):
    c.propagate
    if isLeaf(c,vars)
      parent.exit(leafstatus(c))
    pos = ...           // from vars based on var-select
    for each child: // based on domain-select
      val = ...         // from values of var based on domain-select
      child.post(vars[pos]=val)
      top.init(c,child)
      queue.push(child)

```

Note that, as the `base_search` combinator is a base combinator, its `exit` message is immaterial (there is no child heuristic of `base_search` that could ever call it). The `start` and `init` messages are empty. Many variants on and generalizations of the above implementation are possible.

**Depth-first search engine.** The engine `dfs` serves as a pseudo-combinator at the **top** of a combinator expression `heuristic` and serves as the `heuristic`'s immediate parent as well. It maintains the **queue** of nodes, a stack in this case. The search starts from a given root node by starting the `heuristic` with that node and then entering it. Each time a node has been processed, new nodes may have been pushed onto the queue. These are popped and entered successively.

```

combinator dfs(heuristic)
  start(root):
    top=this
    heuristic.parent=this
    queue=new stack()
    heuristic.start(root)
    heuristic.enter(root)
    while not queue.empty
      heuristic.enter(queue.pop())

  init(n,c):
    heuristic.init(n,c)

```

The engine's `exit` message is empty, the `enter` message is never called and the `init` message delegates initialization to the `heuristic`.

Other engines may be formulated with different queuing strategies.

### 3.3 Combinator Composition

The idea of search combinators is to augment a `base_search`. We illustrate this with a very simple print combinator that prints out every solution as it is found. For simplicity we assume

a solution is just a set of constraint variables *vars* that is supplied as a parameter. Hence, we obtain the basic search setup with solution printing with:

```
dfs (print (vars, base_search (vars, strategy)))
```

**Print.** The print combinator is parametrized by a set of variables *vars* and a search combinator *child*. Implicitly, in a composition, that *child*'s *parent* is set to the print instance. The same holds for all following search combinators with one or more children.

The only message of interest for print is `exit`. When the exit status is **success**, the combinator prints the variables and propagates the message to its parent.

```
combinator print (vars, child)
  exit (c, status) :
    if status == success
      print c.vars
      parent.exit (c, status)
```

The other messages are omitted. Their behavior is default: they all propagate to the child. The same holds for the omitted messages of following unary combinators.

### 3.4 Binary Combinators

Binary combinators are one step up from unary ones. They combine two complete search heuristics into a composite one. The most basic binary combinator is the binary version of `and`. For instance, if we need to label two sets of variables, we can do so with

```
and (base_search (vars1, ...), base_search (vars2, ...))
```

The principle shown here easily generalizes to *n*-ary combinators.

**And.** The (binary) `and` combinator has two children, *left* and *right*. In order to keep track of what child combinator is handling a particular node, the `and` combinator associates with every node an `inLeft` Boolean variable. The **local** keyword indicates that every node has its own instance of that variable. We denote the instance of the `inLeft` variable associated with node *c* as *c.inLeft*.

When `entering` a node, it is delegated to the *left* or *right* combinator based on `inLeft`. At the `start`, the root node is delegated to the *left* combinator, so its `inLeft` variable is set to **true**. The value of `inLeft` is inherited in `init` from the current node to its children. Upon a successful `exit` for *left*, the leaf node becomes the root of a new subtree that is further handled by the *right* combinator.

```
combinator and (left, right) {
  local bool inLeft

  start (root) :
    root.inLeft = true
    left.start (root)

  enter (c) :
    if c.inLeft
```

```

    left.enter(c)
  else
    right.enter(c)

exit(c, status):
  if c.inLeft and status==success
    c.inLeft=false
    right.start(c)
    right.enter(c)
  else
    parent.exit(c, status)

init(p, c):
  c.inLeft=p.inLeft
  if c.inLeft
    left.init(p, c)
  else
    right.init(p, c)

```

Note that the `right` combinator is `started` repeatedly, once for each leaf node of `left`. In general, each combinator can be managing multiple subtrees of the search.

Multiple and combinators may be handling a search node at the same time. For instance in a heuristic of the form `and(and( $s_1, s_2$ ),  $s_3$ )`, two `and` combinators are active at the same time. The scoping of the associated variables works in such a way that each `and` has its own instance of `inLeft` for each node.

### 3.5 Reusable Combinators

Now we show how a *monolithic* combinator can be decomposed into more primitive combinators that can be reused for other purposes.

**Monolithic Combinator.** We start from the following `limitsolutions` combinator that prunes the search after `cutoff` solutions have been found. One new concept is the notion of a global variable associated with a (sub)tree: all descendants of `root` (implicitly) share the same instance of `count`. Hence, any update of `count` by one node is seen by all other nodes in the (sub)tree.

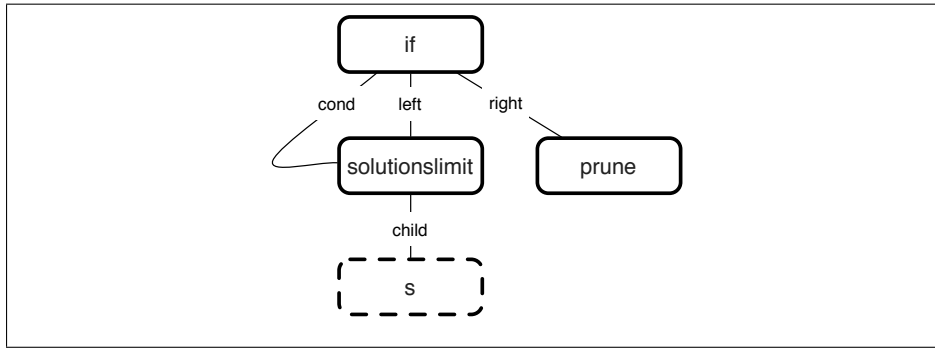
```

combinator limitsolutions(cutoff, child)
  global int count

  start(root):
    root.count = 0
    child.start(root)

  enter(c):
    if count == cutoff
      parent.exit(abort)
    else

```



**Fig. 4:** The decomposition of the limitsolutions combinator

```

    child.enter(c)

exit(c, status) :
    if status==success
        c.count++
    parent.exit(c, status)

```

**Decomposition.** We can split up the above limitsolutions combinator into three different combinators: ifthenelse, solutionslimit and prune. They form a directed acyclic graph as depicted in Figure 4 or denoted as an expression with sharing below:

$$\text{limitsolutions}(\text{cutoff}, s) = \text{ifthenelse}(s', s', \text{prune})$$

where

$$s' = \text{solutionslimit}(\text{cutoff}, s)$$

Here, solutionslimit monitors the number of solutions produced by  $s$ . If this number reaches the `cutoff`, then ifthenelse switches to `prune`, which discards the remaining nodes in the tree.

We now elaborate on each of these combinators individually.

**Prune.** The prune combinator is a minimal base combinator that immediately exits every node with the **abort** status. The start message is empty, and the exit and init messages are never called.

```

combinator prune ()
    enter(c) :
        parent.exit(c, abort)

```

**Solutions Count.** The solutionslimit combinator below illustrates how statistics gathering combinators are implemented. It implements a sub-protocol of **combinator** with an extra message eval that queries the current Boolean value:

```

protocol condition extends combinator
    eval(currentNode);

```



In the case of `solutionslimit`, the returned Boolean value is whether a particular number (`cutoff`) of solutions has not yet been reached by its `child`. For this purpose it maintains the number of solutions found so far in a global variable.

```
condition solutionslimit(cutoff, child)
global int count

start(root):
    root.count = 0
    child.start(root)

exit(c, status):
    if status==success
        c.count++
    parent.exit(c, status)

eval(c):
    return c.count <= cutoff
```

**Ifthenelse.** The `ifthenelse` combinator is parametrized by one `condition` and two child combinators. It associates with every node whether it is handled by the `left child` (`inLeft`); this is the case for the `root` node. Whenever a node `c` is entered that is `inLeft`, the condition is checked. If the condition fails, `c` becomes the root of a subtree that is further handled by `right`.

```
combinator ifthenelse(cond, left, right)
local bool inLeft

start(root):
    root.inLeft=true
    left.start(root)

enter(c):
    if not c.inLeft
        right.enter(c)
    else if cond.eval()
        left.enter(c)
    else
        c.inLeft=false
        right.start(c)
        right.enter(c)

init(p, c):
    c.inLeft=p.inLeft
    if c.inLeft
        left.init(p, c)
    else
        right.init(p, c)
```

### 3.6 Restarting Combinators

Restarting the search is common to several combinators; the mechanic is illustrated below in the `portfolio` combinator.

**Portfolio.** Like the `ifthenelse` and `and` combinators, the `portfolio` combinator switches between child combinators. Only the logic for switching is more complex. In order to simplify presentation, we again restrict the code to the binary case; the  $n$ -ary variant is a straightforward generalization.

Firstly, `portfolio` keeps track of a global “reference” count `ref` of unprocessed nodes to be handled by the `s1` child. This count is incremented whenever a new child node is initialized, and decremented whenever a node is entered for actual processing.

When the last node of `s1` exits (witnessed by the reference count being 0) and the search was not exhaustive, the search starts over from the root, but now with the `s2` child. In order to decide about exhaustiveness, the `portfolio` combinator registers whether any `exit` with status `abort` occurred. At the same time it converts an `abort` inside `s1` into a `failure`, because the `s2` combinator may still perform an exhaustive search and avoid overall non-exhaustiveness. In order to restart from the `root`, a copy of the `root` node is made at the `start`.

Upon a successful `exit`, the leaf node becomes the root of a new subtree that is further handled by the `s2` combinator.

```
combinator portfolio(s1, s2)
  global node copy
  global bool inLeft
  global bool exhaustive
  global int ref

  start(root) :
    copy=root.copy()
    root.inLeft=true
    root.exhaustive=true
    root.ref=1
    s1.start(root)

  enter(c) :
    if c.inLeft
      ref--
      s1.enter(c)
    else
      s2.enter(c)

  exit(c, status) :
    if not c.inLeft
      parent.exit(c, status)
    else
      if status==abort
        status=failure
        c.exhaustive=false
```

```

    if c.ref==0
        if c.exhaustive
            parent.exit(c,status)
        else
            copy.inLeft=false
            s2.start(copy)
            self.enter(copy)
        else
            parent.exit(c,status)

init(p,c):
    ref++;
    if c.inLeft
        s1.init(p,c)
    else
        s2.init(p,c)

```

## 4 Modular Combinator Implementation

The message-based combinator approach lends itself well to different implementation strategies. In the following we briefly discuss two diametrically opposed approaches we have explored: *dynamic composition* (interpretation) and *static composition* (compilation). Using these different approaches, combinators can be adapted to the implementation choices of existing solvers. Sect. 5 shows that both implementation approaches have competitive performance.

### 4.1 Dynamic Composition

To support dynamic composition, we have implemented our combinators as C++ classes whose objects can be allocated and composed into a search specification at runtime. The protocol events correspond to virtual method calls between these objects. For the delegation mechanism from one object to another, we explicitly encode a form of dynamic inheritance called *open recursion* or *mixin inheritance* [4]. In contrast to the OOP inheritance built into C++ and Java, this mixin inheritance provides two essential abilities: 1) to determine the inheritance graph *at runtime* and 2) to use multiple copies of the same combinator class at different points in the inheritance graph. In contrast, C++'s built-in static inheritance provides neither.

The C++ library currently builds on top of the Gecode constraint solver [23]. However, the solver is accessed through a layer of abstraction that is easily adapted to other solvers (e.g., we have a prototype interface to the Gurobi MIP solver). The complete library weighs in at around 2500 lines of code, which is even less than Gecode's native search and branching components.

## 4.2 Static Composition

In a second approach, also on top of Gecode, we statically compile a search specification to a tight C++ loop. Again, every combinator is a separate module independent of other combinator modules. A combinator module now does not directly implement the combinator's behavior. Instead it implements a code generator (in Haskell), which in turn produces the C++ code with the expected behavior.

Hence, our search language compiler parses a search specification, and composes (in mixin-style) the corresponding code generators. Then it runs the composite code generator according to the message protocol. The code generators produce appropriate C++ code fragments for the different messages, which are combined according to the protocol into the monolithic C++ loop. This C++ code is further post-processed by the C++ compiler to yield a highly optimized executable.

As for dynamic composition, the mixin approach is crucial, allowing us to add more combinators without touching the existing ones. At the same time we obtain with the press of a button several 1000 lines of custom low-level code for the composition of just a few combinators. In contrast, the development cost of hand crafted code is prohibitive.

A *compromise* between the above two approaches, itself static, is to employ the built-in mixin mechanism (also called *traits*) available in object-oriented languages such as Scala [6] to compose combinators. A dynamic alternative is to generate the combinator implementations using dynamic compilation techniques, for instance using the LLVM (Low Level Virtual Machine) framework. These options remain to be explored.

## 4.3 MiniZinc with Combinators

As a proof of concept and platform for experiments, we have integrated search combinators into a complete MiniZinc toolchain, comprising a pre-compiler and a FlatZinc interpreter.

The pre-compiler is necessary to support arbitrary expressions in annotations, such as the condition expressions for an `ifthenelse`. The expressions are translated into standard MiniZinc annotations that are understood by the FlatZinc interpreter. User-defined variables have type-`inst svar int` and can be introduced using the standard MiniZinc `let` construct. The `annotation` construct of MiniZinc has been extended to support simple function definitions. The following example shows a MiniZinc version of the restart-based branch-and-bound heuristic from Sect. 2.3:

```

annotation limit(var bool: c, ann: s) =
    ifthenelse(c,s,prune);

annotation once(ann: s) = limit(solutions < 1, s);

annotation rbab(var int: obj, ann: s) =
    let { svar int: best = MAXINT } in
    restart(true, and([
        post(obj < best),
        once(s),
        assign(best,obj)]));

solve ::rbab(x,int_search(y,input_order,assign_lb)) satisfy;
```

The pre-compiler translates this code as follows:

```
solve :: sh_let(sh_letvar("best"), sh_int(MAXINT),
  sh_restart(sh_cond_true, sh_and([
    sh_post_succeed(sh_cond_lt(sh_intvar(objective),
      sh_letvar("best"))),
    sh_let(sh_letvar("solutioncount"), 0,
      sh_ifthenelse(sh_cond_lt(sh_letvar("solutioncount"),
        sh_int(1)),
        sh_solutioncount(sh_letvar("solutioncount"),
          sh_int_search(x, sh_var_input_order,
            sh_val_assign_lb)),
        sh_prune))),
    sh_assign(sh_letvar("best"), sh_intvar(objective))))))
satisfy;
```

All literals are quoted (e.g. `sh_int(1)`), user-defined search variables are turned into quoted strings (`lv("best")`), expressions like `obj < best` are translated into annotation terms (`sh_cond_lt ...`), and statistics are made explicit, introducing search variables and special combinators (`sh_solutioncount`). The result of the pre-compilation is valid, well-typed MiniZinc, which is then passed through the standard `mzn2fzn` translator to produce FlatZinc ready for solving. We intend to incorporate the translations done by the pre-compiler into the standard `mzn2fzn` in the future.

We extended the Gecode FlatZinc interpreter to parse the search combinator annotation and construct the corresponding heuristic using the Dynamic Composition approach described above. The three tools, pre-compiler, `mzn2fzn`, and the modified FlatZinc interpreter thus form a complete toolchain for solving MiniZinc models using search combinators. The source code including examples can be downloaded from

<http://www.gecode.org/flatzinc.html>.

## 5 Experiments

This section evaluates the performance of our two implementations. It establishes that a search heuristic specified using combinators is competitive with a custom implementation of the same heuristic, exploring exactly the same tree.

Sect. 3.1 introduced a message protocol that defines the communication between the different combinators *for one node of the search tree*. Any overhead of a combinator-based implementation must therefore come from the processing of each node using this protocol. All combinators discussed earlier process each message of the protocol in constant time (except for the `base_search` combinators, of course). Hence, we expect at most a constant overhead per node compared to a native implementation of the heuristic.

In the following, two sets of experiments confirm this expectation. The first set consists of artificial benchmarks designed to expose the overhead per node. The second set consists of realistic combinatorial problems with complex search strategies.

The experiments were run on a 2.26 GHz Intel Core 2 Duo running Mac OS X. The results are the averages of 10 runs, with a coefficient of deviation less than 1.5%.

**Stress Test.** The first set of experiments measures the overhead of calling a single combinator during search. We ran a complete search of a tree generated by 7 variables with domain

$\{0, \dots, 6\}$  and *no* constraints (1 647 085 nodes). To measure the overhead, we constructed a basic search heuristic *s* and a stack of *n* combinators:

`portfolio([portfolio([...portfolio([s,prune])...],prune)],prune])`

where *n* ranges from 0 to 20 (realistic combinator stacks, such as those from the examples in this article, are usually not deeper than 10). The numbers in the following table report the runtime with respect to using the plain heuristic *s*, for both the static and the dynamic approach:

<i>n</i>	1	2	5	10	20
<i>static %</i>	106.6	107.7	112.0	148.3	157.5
<i>dynamic %</i>	107.3	117.6	145.2	192.6	260.9

A single combinator generates an overhead of around 7%, and 10 combinators add 50% for the static and 90% for the dynamic approach. In absolute runtime, however, this translates to an overhead of around 17 ms (70 ms) per million nodes and combinator for the static (dynamic) approach. Note that this is a worst-case experiment, since there is no constraint propagation and almost all the time is spent in the combinators.

**Benchmarks.** The second set of experiments shows that in practice, this overhead is dwarfed by the cost of constraint propagation and backtracking. Note that the experiments are not supposed to demonstrate the best possible search heuristics for the given problems, but that a search heuristic implemented using combinators is just as efficient as a native implementation.

Fig. 5 compares Gecode’s optimization search engines with branch-and-bound implemented using combinators. On the well-known Golomb Rulers problem, both dynamic combinators and native Gecode are slightly slower than static combinators. Native Gecode uses dynamically combined search heuristics, but is much less expressive. That is why the static approach with its specialization yields better results.

On the radiotherapy problem (see Sect. 2.3), the dynamic combinators show an overhead of 6–11%. For native Gecode, `exh_once` must be implemented as a nested search, which performs similarly to the dynamic combinators. However, in instances 5 and 6, the compiled combinators lose their advantage over native Gecode. This is due to the processing of `exh_once`: As soon as it is finished, the combinator approach processes all nodes of the `exh_once` tree that are still in the search queue, which are now pruned by `exh_once`. The native Gecode implementation simply discards the tree. We will investigate how to incorporate this optimization into the combinator approach.

The job shop scheduling examples, using the combination of branch-and-bound and discrepancy limit discussed in Sect. 2.3, show similar behavior. In ABZ1-5 and mt10, the interpreted combinators show much less overhead than in the short-running instances. This is due to more expensive propagation and backtracking in these instances, reducing the relative overhead of executing the combinators.

In summary, the experiments show that the expressiveness and flexibility of a rich combinator-based search language can be achieved without any runtime overhead in the case of the static approach, and little overhead for the dynamic version.

## 6 Related Work

This section explores and discusses previous work that is closely related to search combinators as presented in this article.

	<i>Compiled</i>	<i>Interpreted</i>	<i>Gecode</i>
<i>Golomb 10</i>	0.61 s	101.8%	102.5%
<i>Golomb 11</i>	12.72 s	102.9%	101.8%
<i>Golomb 12</i>	125.40 s	100.6%	101.9%
<i>Radiotherapy 1</i>	71.13 s	105.9%	107.3%
<i>Radiotherapy 2</i>	6.22 s	110.9%	110.1%
<i>Radiotherapy 3</i>	11.78 s	108.3%	108.1%
<i>Radiotherapy 4</i>	16.44 s	107.5%	106.9%
<i>Radiotherapy 5</i>	69.89 s	108.1%	98.7%
<i>Radiotherapy 6</i>	106.04 s	109.2%	99.1%
<i>Job Shop G2</i>	7.25 s	146.3%	101.2%
<i>Job-Shop G4</i>	6.96 s	164.0%	107.75%
<i>Job-Shop H1</i>	38.05 s	153.1%	103.81%
<i>Job Shop H3</i>	52.02 s	162.5%	102.8%
<i>Job Shop H5</i>	20.88 s	153.2%	107.0%
<i>Job Shop ABZ1-5</i>	2319.00 s	103.7%	100.1%
<i>Job Shop mt10</i>	2181.00 s	104.5%	99.9%

Fig. 5: Experimental results

## 6.1 MCP

This work directly extends our earlier work on **Monadic Constraint Programming** (MCP) [20]. MCP introduces stackable search transformers, which are a simple form of search combinators, but only provide a much more limited and low level form of search control. In trying to overcome its limitations we arrived at search combinators.

## 6.2 Constraint Logic Programming

Constraint logic programming languages such as **ECLiPSe** [7] and **SICStus Prolog** [25] provide programmable search via the built-in search of the paradigm, allowing the user to define *goals* in terms of conjunctive or disjunctive sub-goals.

Prolog’s limitation is that it does not permit cross-cutting between goals. For instance, disjunctions inside goals are too well encapsulated to observe them or interfere with them from outside that goal. Hence, combinators that inject additional behavior in disjunctions, i.e. to observe and/or prune the number of branches, cannot be expressed in a modular way. In contrast, cross-cutting is a crucial feature of our combinator approach, where a combinator higher up in the stack can interfere with a sub-combinator, while remaining fully compositional. In summary, apart from conjunction and disjunction, Prolog’s goal-based heuristics cannot be combined arbitrarily.

ECLiPSe copes with this limitation by combining a limited number of search heuristics into a monolithic `search/6` predicate. With various parameters the user controls which of the heuristics is enabled (e.g., depth-bounded, node-bounded or limited discrepancy search). A fixed number of compositions are supported, such as changing strategy when the depth bound finishes. The labeling itself is user programmable. If a user is not happy with the set of supported heuristics in `search/6`, he has to program his own from scratch.

**IBM ILOG CP Optimizer** [5] supports Prolog-style goals in C++ [16], and like Prolog goals, these do not support cross-cutting.

### 6.3 The Comet Language

The **Comet** [27] system features fully programmable search [28], built upon the basic concept of *continuations*, which make it easy to capture the state of the solver and write non-deterministic code.

The Comet library provides abstractions like the non-deterministic primitives `try` and `tryall` that split the search specification in two (orthogonal) parts: 1) the specification of the search tree which corresponds to our `base_search` heuristics, and 2) the exploration of that search tree by means of a *search controller*. In terms of our approach, the search controller determines both the queueing strategy and the behavior of the search heuristic (minus the base search) within a single entity. In other words, it defines what to do when starting or ending a search, failing, or adding a new choice.

Complex heuristics can be constructed as custom controllers, either by inheriting from existing controllers or implementing them from scratch.

Albeit at a different level of abstraction (e.g., compare the Comet definition of depth-bounded search in Figure 6 to the combinator definition  $\text{dbs}(n, s) \equiv \text{limit}(\text{depth} \leq n, s)$ ), search controllers are quite similar to combinators as presented in this article. However, there is one essential difference. Our combinators are meant to be compositional, whereas search controllers are not. This difference in spirit is clearly reflected in 1) the design of the interface and its associated protocol, and 2) the instances:

1. The design of search controllers is simpler than that of search combinators because it does not take compositionality into account. While many of the messages in the two approaches are similar in spirit, the search combinator approach also stipulates the flow of messages within a search combinator composition. Notably, while most of the messages propagate top-down through a combinator stack, it is vital to compositionality that the `exit` message proceeds in a bottom-up manner. For instance, this bottom-up flow enables the inner and combinator in the composition  $\text{and}(\text{and}(s_1, s_2), s_3)$  to intercept leaf nodes of  $s_1$  and start  $s_2$  before its parent starts  $s_3$ . The other way around would clearly exhibit an undesirable semantics. In Comet, this compositional protocol is entirely absent. All messages are directed at the single search controller.
2. In terms of instantiation, because of their compositional nature, we promote many “small” combinator instances that each capture a single primitive feature. This approach provides us with a high-level modeling language for search, as the primitive combinators are conveniently assembled into many different search heuristics. In contrast, all Comet search controller instances we are aware of<sup>2</sup> are essentially monolithic implementations of a particular search heuristic; none of them takes other search controllers as arguments. Through a common abstract base class the instances share some basic infrastructure, but to implement a new search controller one basically starts from scratch.

The fact that search controllers have not been designed with compositionality in mind obviously does not mean that compositionality cannot be achieved in Comet. On the contrary, we believe that it is most easily achieved by integrating search controllers with the compositional design of our search combinators. In fact, because of Comet’s powerful primitives for non-determinism, this would lead to a particularly elegant implementation.

<sup>2</sup> i.e., those published in papers and shipped with the Comet library.



```

class DBS extends AbstractSearchController {
  stack{Continuation} s;
  int limit;
  DBS(SearchSolver so, int n) : AbstractSearchController(so) {
    s = new stack{Continuation}();
    limit = n;
  }
  void startTry() {
    if (s.getSize() > limit) fail();
  }
  void addChoice(Continuation f) {
    s.push(f);
  }
  void fail() {
    if (s.empty())
      exit();
    else
      call(s.pop());
  }
}

```

Fig. 6: Definition of depth-bounded search in Comet.

#### 6.4 Other Systems

The **Salsa** [12] language is an imperative domain-specific language for implementing search algorithms on top of constraint solvers. Its center of focus is a node in the search process. Programmers can write custom *Choice* strategies for generating next nodes from the current one; Salsa provides a regular-expression-like language for combining these Choices into more complex ones. In addition, Salsa can run custom procedures at the *exits* of each node, right after visiting it. We believe that Salsa’s Choice construct is orthogonal to our approach and could be incorporated. Custom exit procedures show similarity to combinators, but no support is provided for arbitrary composition.

**Oz** [26] was the first language to truly separate the definition of the constraint model from the exploration strategy [22]. Computation spaces capture the solver state and the possible choices. Strategies such as DFS, BFS, LDS, Branch and Bound and Best First Search are implemented by a combination of *copying* and *recomputation* of computation spaces. The strategies are monolithic, there is no notion of search combinators.

**Zinc/MiniZinc** [13, 14] lets the user specify search in its *annotation language*. There is a proposal for a more expressive search language for MiniZinc [19], but it is limited to basic variable ordering and domain splitting strategies. For Zinc, a language extension is available for implementing variable selection and domain splitting [17] but again it does not address more than basic search.

#### 6.5 Autonomous Search

Autonomous search (AS) [8] addresses the challenge of providing complex application-tailored search heuristics in a different way. Rather than leaving the specification and tuning of search heuristics to the programmer, AS promotes systems that autonomously self-tune their performance while solving problems. Hence, while search combinators make writing

search heuristics easier, AS takes it out of the hands of the programmer altogether. Well-known instances of this approach are Impact Based Search [18] or the *weighted degree* heuristic [3].

AS has advantages for 1) smaller problems where it produces a decent heuristic without programmer investment, and for 2) novice users who don't know how to obtain a decent heuristic. However, loss of programmer control is a liability for hard problems where AS can be ineffective and often only expert knowledge makes the difference.

## 7 Conclusion

We have shown how combinators provide a powerful high-level language for modeling complex search heuristics. To make this approach useful in practice, we devised an architecture in which the modularity of the language is matched by the modularity of the implementation. This relieves system developers from a high implementation cost and yet, as our experiments show, imposes no runtime penalty.

For future work, parallel search on multi-core hardware fits perfectly in our combinator framework. We have already performed a number of preliminary experiments and will further explore the benefits of search combinators in a parallel setting. We will also explore potential optimizations (such as the short-circuit of `exh_once` from Sect. 5) and different compilation strategies (e.g., combining the static and dynamic approaches from Sect. 4).

In addition we consider to apply search combinators in other problem domains like Mixed Integer Programming (MIP) and  $A^*$  where search strategies have a major impact on performance and no dominant default search exists. Finally, combinators need not necessarily be heuristics that control the search. They may also monitor search, e.g., by gathering statistics or visualizing the search tree.

**Acknowledgements** NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council. This work was partially supported by Asian Office of Aerospace Research and Development grant 10-4123.

## References

1. Baatar, D., Boland, N., Brand, S., Stuckey, P.J.: CP and IP approaches to cancer radiotherapy delivery optimization. *Constraints* **16**(2), 173–194 (2011)
2. Boussemart, F., Hemery, F., Lecoutre, C., Sais, L.: Boosting systematic search by weighting constraints. In: ECAI'04, pp. 146–150 (2004)
3. Boussemart, F., Hemery, F., Lecoutre, C., Sais, L.: Boosting systematic search by weighting constraints. In: R.L. de Mántaras, L. Saitta (eds.) ECAI, pp. 146–150. IOS Press (2004)
4. Cook, W.R.: A denotational semantics of inheritance. Ph.D. thesis, Brown University (1989)
5. IBM ILOG CP Optimizer (2011). URL <http://www-01.ibm.com/software/integration/optimization/cplex-cp-optimizer/>
6. Cremet, V., Garillot, F., Lenglet, S., Odersky, M.: A core calculus for Scala type checking. In: R. Kralovic, P. Urzyczyn (eds.) MFCS, LNCS, vol. 4162, pp. 1–23. Springer, Heidelberg (2006)
7. ECLiPSe. <http://www.eclipse-clp.org/> (2008)
8. Hamadi, Y., Monfroy, E., Saubion, F. (eds.): *Autonomous Search*. Springer-Verlag (2012)
9. Harvey, W.D., Ginsberg, M.L.: Limited discrepancy search. In: IJCAI, pp. 607–613 (1995)
10. Kiczales, G., Lamping, J., Menhdhekar, A., Maeda, C., Lopes, C., Loingtier, J., Irwin, J.: Aspect-oriented programming. In: ECOOP'97, pp. 220–242 (1997)
11. Korf, R.E.: Depth-first iterative-deepening: an optimal admissible tree search. *Artif. Intell.* **27**, 97–109 (1985)
12. Laburthe, F., Caseau, Y.: SALSAs: A language for search algorithms. *Constraints* **7**(3), 255–288 (2002)

13. Marriott, K., Nethercote, N., Rafeh, R., Stuckey, P., Garcia de la Banda, M., Wallace, M.: The design of the Zinc modelling language. *Constraints* **13**(3), 229–267 (2008)
14. Nethercote, N., Stuckey, P.J., Becket, R., Brand, S., Duck, G.J., Tack, G.: MiniZinc: Towards a standard CP modelling language. In: C. Bessiere (ed.) *CP, LNCS*, vol. 4741, pp. 529–543. Springer, Heidelberg (2007)
15. Perron, L.: Search procedures and parallelism in constraint programming. In: J. Jaffar (ed.) *CP, LNCS*, vol. 1713, pp. 346–360. Springer, Heidelberg (1999)
16. Puget, J.F.: A C++ implementation of CLP. In: Proceedings of the Second Singapore International Conference on Intelligent Systems (SPICIS), pp. B256–B261 (1994)
17. Rafeh, R., Marriott, K., de la Banda, M.G., Nethercote, N., Wallace, M.: Adding search to zinc. In: P.J. Stuckey (ed.) *CP, LNCS*, vol. 5202, pp. 624–629. Springer (2008)
18. Refalo, P.: Impact-based search strategies for constraint programming. In: M. Wallace (ed.) *CP, LNCS*, vol. 3258, pp. 557–571. Springer, Heidelberg (2004)
19. Samulowitz, H., Tack, G., Fischer, J., Wallace, M., Stuckey, P.: Towards a lightweight standard search language. In: *ModRef* (2010)
20. Schrijvers, T., Stuckey, P.J., Wadler, P.: Monadic constraint programming. *Journal of Functional Programming* **19**(6), 663–697 (2009)
21. Schrijvers, T., Tack, G., Wuille, P., Samulowitz, H., Stuckey, P.J.: Search combinators. In: J.H. Lee (ed.) *CP’11, LNCS*, vol. 6876, pp. 774–788. Springer (2011)
22. Schulte, C.: Programming constraint inference engines. In: G. Smolka (ed.) *CP, LNCS*, vol. 1330, pp. 519–533. Springer, Heidelberg (1997)
23. Schulte, C., et al.: Gecode, the generic constraint development environment (2009). <http://www.gecode.org/>
24. Sellmann, M., Kadioglu, S.: Dichotomic search protocols for constrained optimization. In: P.J. Stuckey (ed.) *CP, LNCS*, vol. 5202, pp. 251–265. Springer, Heidelberg (2008)
25. SICStus Prolog. <http://www.sics.se/isl/sicstuswww/site/> (2008)
26. Smolka, G.: The Oz programming model. In: J. van Leeuwen (ed.) *Computer Science Today, LNCS*, vol. 1000, pp. 324–343. Springer, Heidelberg (1995)
27. Van Hentenryck, P., Michel, L.: *Constraint-Based Local Search*. MIT Press (2005)
28. Van Hentenryck, P., Michel, L.: Nondeterministic control for hybrid search. *Constraints* **11**(4), 353–373 (2006)
29. Van Hentenryck, P., Perron, L., Puget, J.F.: Search and strategies in OPL. *ACM TOCL* **1**(2), 285–315 (2000)